NASA Technical Memorandum 107489
AIAA–97–3295

Army Research Laboratory
Memorandum Report ARL–MR–369

Parallel ALLSPD–3D: Speeding Up Combustor Analysis Via Parallel Processing

David M. Fricker
U.S. Army Research Laboratory
Lewis Research Center
Cleveland, Ohio

Prepared for the

PARALLEL ALLSPD-3D: SPEEDING UP COMBUSTOR ANALYSIS VIA PARALLEL PROCESSING

David M. Fricker*
U.S. Army Research Laboratory - Vehicle Technology Center, Lewis Site
Cleveland, OH

Abstract

The ALLSPD-3D Computational Fluid Dynamics code for reacting flow simulation was run on a set of benchmark test cases to determine its parallel efficiency. These test cases included non-reacting and reacting flow simulations with varying numbers of processors. Also, the tests explored the effects of scaling the simulation with the number of processors in addition to distributing a constant size problem over an increasing number of processors. The test cases were run on a cluster of IBM RS/6000 Model 590 workstations with ethernet and ATM networking plus a shared memory SGI Power Challenge L workstation. The results indicate that the network capabilities significantly influence the parallel efficiency, i.e., a shared memory machine is fastest and ATM networking provides acceptable performance. The limitations of ethernet greatly hamper the rapid calculation of flows using ALLSPD-3D.

Nomenclature

S = Speedup
E = Efficiency
N = Number of processors
T = Time
T_{wall} = wall clock or elapsed time
T_{cpu} = CPU time used by process
$_{serial}$ = serial processing with a single processor
$_{parallel}$ = parallel processing with multiple processors
ATM = Asynchronous Transfer Mode network
ethernet = Ethernet network
$Re_{d,a}$ = Reynolds Number based on diameter
T_{ref} = Reference Temperature
U_{ref} = Reference Velocity
K = Kelvin
m/s = meters/second

Introduction

ALLSPD-3D Capabilities

The ALLSPD-3D combustion code is a numerical tool developed by the Internal Fluid Mechanics Division (which is now the Turbomachinery and Propulsion Systems Division) at the NASA Lewis Research Center for simulating chemically reacting flows in aerospace propulsion systems.[1] It provides the designer of advanced engines an analysis tool that employs state-of-the-art computational technology. The code can simulate multi-phase, swirling flows over a wide Mach-number range in combustors of complex geometry. Three-dimensional, curvilinear, structured grids with multiple zones and internal obstacles give great flexibility in fitting the grid to solid bodies in the flow simulation. Various boundary conditions (multiple inlets/outlets, dilution holes, transpiration holes, periodic, symmetry, far-field, adiabatic or isothermal walls, centerline singularity) also increase the utility of ALLSPD-3D in solving complex flow simulations.

The ALLSPD-3D Computational Fluid Dynamics (CFD) code which was released in November, 1995, evolved from the two-dimensional code ALLSPD-2D (released in June, 1993). Besides extension to three dimensions, the newer code featured several improvements and enhancements, including a user-friendly Graphical User Interface (GUI), multi-platform capability (supercomputers, workstations, and parallel processors), improved turbulence and spray models, and more generalized property and chemical reactions databases. Also, eddy breakup models for turbulence-chemistry interactions were introduced. A very warmly received feature of the ALLSPD-3D version 1.0 code was the GUI for easier problem setup and post-processing.

*Engine Components Division.
Senior Member AIAA.

The ALLSPD combustion codes utilize a finite-difference, compressible flow formulation with low Mach number preconditioning of the Navier-Stokes equations. (The ALLSPD-3D code is intended only for subsonic flow simulations since it uses central-differencing for convective and viscous terms on right and left-hand sides.) Laminar or turbulent flow capability also exists, and the turbulent flows are solved using a low-Reynolds number k-ε turbulence model. The chemistry model can handle frozen or finite rate chemistry flows. Spray combustion is supported by a stochastic, separated flow spray model.

Need for parallelization

ALLSPD was parallelized in response to the changing computational capabilities of the major engine companies, specifically, the move from large supercomputers to small workstations. ALLSPD-3D is memory and CPU intensive for practical engineering problems. This led to the need for parallel processing on UNIX workstations such as those from HP, IBM, SGI, & Sun. However, the serial code was not to be abandoned, nor was the parallel version to be wildly divergent from the serial code. Also, the parallel code needed to be developed using parallel processing techniques readily available to the average user. Therefore, ALLSPD-3D was parallelized using the de-facto standard PVM (Parallel Virtual Machine) message passing library and with minimal modifications to the serial code.

Transferring data by message passing supplies exactly the information a process needs from its neighboring zones without requiring memory space for all of the data in all of the other zones. Because each process needs data for only its own grid zone (including those ghost cells which actually belong to neighboring zones), each process only needs enough memory for the largest zone. This reduced memory feature of parallel processing can be very beneficial with large problem sizes. Also, since each process only calculates data on its zone, the time needed to calculate a single iteration is reduced to approximately the time needed for the most numerically intensive zone. The only cost for these great benefits of parallel processing is the time it takes to transfer data between neighbors.

ALLSPD-3D Parallelization

Domain decomposition

The parallel processing in ALLSPD-3D is quite simple: the code is inherently divided in the data domain, therefore domain decomposition is used. The multiple grid zone feature provides natural dividing lines in the data for decomposing the problem onto multiple processors, i.e., each grid zone is a natural candidate for parallel processing. This also minimizes the changes to the serial code. Boundary data is exchanged between processors using the PVM message-passing library, and each processor only needs as much memory as demanded by the largest grid zone. This memory limitation is due to the lack of dynamic memory allocation in ALLSPD-3D; all array sizes are set at compile time based upon the largest grid zone since it falls within the Single Program, Multiple Data (SPMD) paradigm. SPMD can be translated as each processor running the same program as all of the other processors but with differing data.

Unfortunately, this limitation extends to the amount of data transferred between processors at the end of each iteration. The first release of ALLSPD-3D contains a design flaw which sets the amount of data to transfer using the maximum possible size of a grid zone's face regardless of how much smaller the grid face being transferred is. The maximum face size is determined at compile time, and this sets the amount of data transferred for all processors. If the size of a particular grid face to be passed to a neighboring grid zone is much smaller than the maximum possible, then a substantial penalty in communication time is taken by the transfer of unneeded information. Reducing this penalty requires code modifications to properly size the amount of data to transfer.

Message passing and PVM

The PVM (Parallel Virtual Machine) message-passing library was developed at Oak Ridge National Laboratory in Oak Ridge, Tennessee.[3] PVM was chosen because of its wide acceptance, installed user base, and portability. PVM is used in a wide variety of applications on numerous architectures and has become a de-facto standard for message-passing libraries.

The PVM library has many features including spawning of processes on a virtual machine and the

communication of various message types between architectures which may have inherently different data structures. These features are used in the parallel version of ALLSPD-3D.

ALLSPD-3D version 1.0b with a minor modification was used for this study of parallel efficiency. The modification involves changing the method used to

was modified for each variation. For simple speedup testing, the baseline grid was split into multiple zones of equal size with one zone per processor. To test the effects of scaling the problem with the number of processors, the baseline grid was mirrored across symmetry planes for the two and four processor cases. Then the four processor grid was refined and divided to create the eight and sixteen processor test

NUMBER OF ZONES	ZONE DIMENSIONS	POINTS PER ZONE	TOTAL NUMBER OF POINTS
1	41 x 21 x 61	52521	52521
2	41 x 21 x 61	52521	105042
4	41 x 21 x 61	52521	210084
8	41 x 21 x 61	52521	420168
16	41 x 21 x 61	52521	840336

Table 3 - Transition duct grids for scaled speedup

Tables 5 and 6 detail the grids used in each transition duct test case.

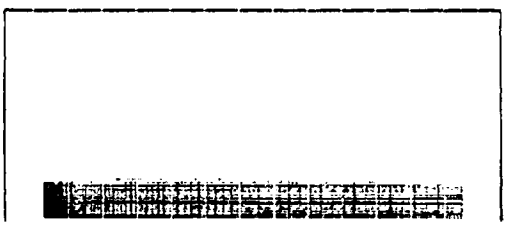

Results

Speedup is defined as the CPU time of the serial code for a particular test case divided by the wall clock or elapsed time of the parallel code for the same test case. The parallel efficiency is the speedup divided by the number of processors.[3] Equations 1 and 2 show these definitions in a more mathematical form.

$$S = \frac{T_{cpu\,serial}}{T_{wall\,parallel}}$$

Equation 1 - Definition of Parallel Speedup

$$E = \frac{S}{N}$$

Equation 2 - Definition of Parallel Efficiency

All test cases were run on dedicated workstations. A cluster of sixteen IBM RS/6000 Model 590 workstations with ethernet and ATM networks and a single SGI Power Challenge L workstation with eight CPUs were used for the tests. The sixteen zone test cases were not run on the SGI Power Challenge L to keep the ratio of one grid zone per processor for all tests. The RS/6000 workstations used PVM version 3.3.10 while the SGI workstation used SGI Array version 2.0 which contains a version of PVM tuned for SGI workstations by SGI.

Each test case was run for 100 iterations and timed with the UNIX command *timex*. This number was normalized using the single processor serial code memory requirement.

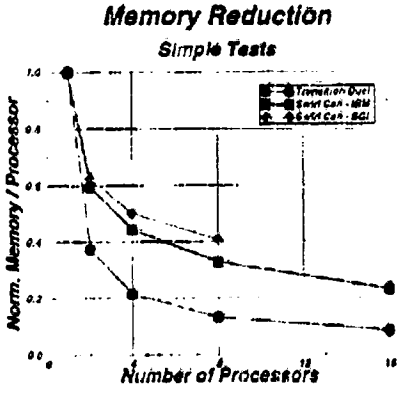

Figure 3

The transition duct shows the most dramatic memory reduction. With four processors, the per processor memory is only about 20% of the single zone test case. Thus, four workstations in parallel would need less aggregate memory than a single machine computing the problem serially because of the way ALLSPD-3D does memory management. Sixteen processors would need less than 10% of the memory needed by the single zone test case on a single CPU workstation. The swirl can test case does not show as

even though sixteen processors are communicating at the same time on every iteration.

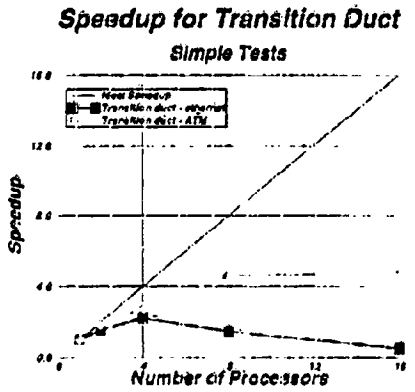

Figure 4

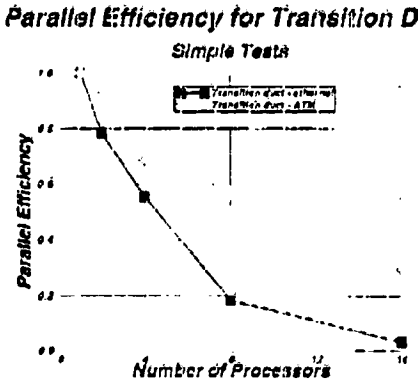

Figure 5

The parallel efficiency for these tests are plotted in Figure 5. Ideal parallel efficiency is 1.0 or 100%, i.e., two processors run twice as fast as one for the same problem. Again, the poor performance of the ethernet network shows itself. ATM networking does encounter a significant drop in parallel efficiency for sixteen processors, but the roughly 60% efficiency with only eight processors is quite acceptable.

The parallel speedup for the swirl can test cases are shown in Figure 6. In addition to the effects of networking on the speedup, we can see the effects of adding chemical reactions and spray modelling to the flow simulation. Adding these features increases the computation to communication ratio for the processors and can also cause the processors to communicate their per iteration results at slightly different times. This would help to reduce the network contention, especially for shared medium networks such as ethernet.

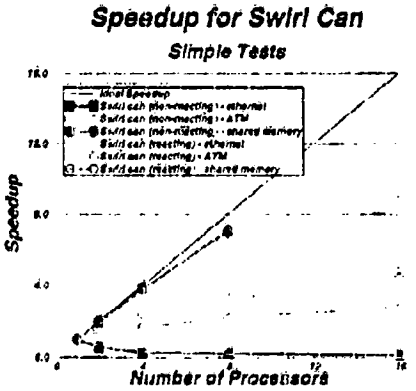

Figure 6

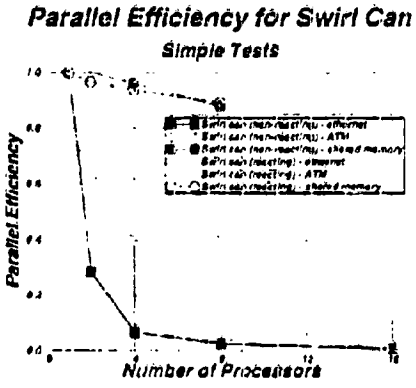

Figure 7

Again, the ethernet test runs show disappointing parallel speedup. This time, however, the ethernet is so overwhelmed by the large data transfer packets hitting the network at the same time that the serial code performs better for all cases. This is because the size of the data packets transferred after every iteration are sized on the maximum possible face. In this case, the actual amount of needed information is much smaller since the zone interfaces are J-K faces and the packets are sized by the J-K faces. The ATM network is decidedly better than the ethernet merely by having speedup values greater than one, but a maximum parallel speedup of only three or four

for the sixteen processor tests is a moot improvement. The shared memory test runs on the SGI Power Challenge L workstation achieve near ideal parallel speedup. As a matter of fact, the two processor test case reaches super-linear speedup. This is most likely due to memory cache effects. In all networks the addition of chemical reactions improves the parallel speedup with the ATM network benefitting most. The shared memory run benefits least from the increase in computation to communication ratio because the shared memory "network" provides

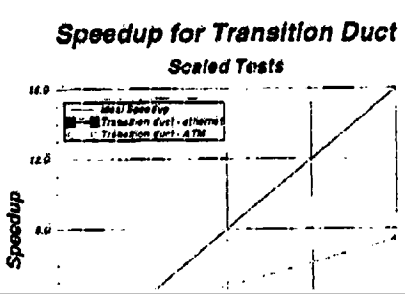

scaling except that the single workstation needs a larger amount of total memory. For all versions, the additional computational burden of chemical reactions has a constant but negligible improvement in parallel speedup.

Concluding Remarks

ALLSPD-3D can simulate flows on clusters of UNIX workstations or multiple processor workstations with shared memory using PVM for data transfer. This

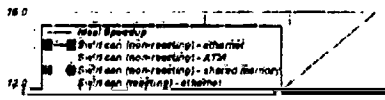

machines, but results in a communication-bound problem with limits on speedup. Faster networks alleviate the situation, but not completely. Shared memory machines provide the fastest communications but can be expensive and require enough memory for the entire problem to be solved.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	June 1997	Technical Memorandum

4. TITLE AND SUBTITLE

Parallel ALLSPD 3D: Speeding Up Combustor Analysis Via Parallel Processing

5. FUNDING NUMBERS

WU-523-26-33

REPRODUCTION QUALITY NOTICE

We use state-of-the-art high speed document scanning and reproduction equipment. In addition, we employ stringent quality control techniques at each stage of the scanning and reproduction process to ensure that our document reproduction is as true to the original as current scanning and reproduction technology allows. However, the following original document conditions may adversely affect Computer Output Microfiche (COM) and/or print reproduction:

- Pages smaller or larger than 8.5 inches x 11 inches.
- Pages with background color or light colored printing.
- Pages with smaller than 8 point type or poor printing.
- Pages with continuous tone material or color photographs.
- Very old material printed on poor quality or deteriorating paper.

If you are dissatisfied with the reproduction quality of any document that we provide, particularly those not exhibiting any of the above conditions, please feel free to contact our Directorate of User Services at (703) 767-9066/9068 or DSN 427-9066/9068 for refund or replacement.

END SCANNED DOCUMENT